

O

AR-009-442

DSTO-TR-0279

T

Application of Artificial Intelligence
in Decision Making in Mine
Countermeasures

T.M. Mansell, D.R. Skinner
and K.K. Benke

S

19960429 012

D

APPROVED FOR PUBLIC RELEASE

DISTRIBUTION STATEMENT A

Approved for public release
Distribution Unlimited

© Commonwealth of Australia

DETC QUALITY INSPECTED 1

DEPARTMENT OF DEFENCE
DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

Application of Artificial Intelligence to Decision Making in Mine Countermeasures

T.M. Mansell, D.R. Skinner and K.K. Benke

**Maritime Operations Division
Aeronautical and Maritime Research Laboratory**

DSTO-TR-0279

ABSTRACT

A number of analytic techniques used in Artificial Intelligence are examined in the context of decision making in mine countermeasures. Attention is directed at five major techniques, involving statistical inference, probabilistic inference, evidential reasoning, fuzzy logic and artificial neural networks. In the cases of statistical inference and evidential reasoning, solutions to appropriate problems are described. Eleven other techniques are dealt with more briefly, in most cases with worked examples of appropriate naval application.

The main conclusion reached is that, in view of the probable shortage of accurate information under operational conditions, evidential reasoning and fuzzy logic are likely to be the most appropriate means for presenting relevant data to decision makers, and that artificial neural networks will be useful for representing complicated or empirical relationships between observed factors.

RELEASE LIMITATION

Approved for public release

DEPARTMENT OF DEFENCE

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

Published by

*DSTO Aeronautical and Maritime Research Laboratory
PO Box 4331
Melbourne Victoria 3001*

*Telephone: (03) 626 8111
Fax: (03) 626 8999
© Commonwealth of Australia 1995
AR No. 009-442
December 1995*

APPROVED FOR PUBLIC RELEASE

Application of Artificial Intelligence to Decision Making in Mine Countermeasures

Executive Summary

The discipline of Artificial Intelligence (AI) is characterised by the development of computational models emulating various aspects of human intelligence. Computer-based AI techniques have possible application whenever the speed and volume of information processing threatens to overwhelm the human resources available. The AI approach is characterised by an accent on symbolic representations and inference rather than being restricted to classical quantitative approaches used in electronic data processing. Very little knowledge in the world is precise, certain, or complete and AI techniques offer a means of processing this uncertain or incomplete information.


In this report, selected AI techniques are investigated in the context of minewarfare modelling and mine countermeasures. Even when the nation or organisation responsible for a mine field can be identified, there may be uncertainty as to the type of mine laid. In addition, any given modern mine can be configured in many ways, with variations in parameters such as the ship count, sensor settings, and the mine-actuation algorithm. The number and location of mines may never be known with any certainty. All of these issues are at present addressed using probabilistic and statistical methods. At an operational level, it is unusual for initial estimates of critical factors to be updated continually on the basis of events that have been experienced, such as the number of mines activated during sweeping. If an unexpected event occurs, operations may be stopped whilst revised tactics are considered. AI methods offer scope for the incorporation of decision-making that is adaptive and partially autonomous.

This investigation revealed that the approaches with most potential for applications in mine countermeasures include evidential reasoning, fuzzy logic and, to a lesser extent, artificial neural networks. Evidential reasoning is a basis for representing uncertain and incomplete information, and provides working tools for manipulating bodies of available evidence. Fuzzy logic deals with a different type of uncertainty to that associated with evidential reasoning - the uncertainty is with respect to the *quantitative* values of factors, rather than in the *confidence* placed on specified conditions. Fuzzy logic can also deal with relationships represented in vaguely defined concepts. Artificial neural networks are best suited to classification problems in domains where good training data are available, and are often associated with domain interpolation rather than extrapolation.

Author

T.M. Mansell

Maritime Operations Division



Todd Mansell received a BSc degree in Physics and Electronics from Deakin University in 1988. He then completed a first class Honours degree with a research project in digital image processing. In 1989 Dr Mansell joined the DSTO as a Cadet Research Scientist, and in 1994 completed a PhD in artificial intelligence at The Melbourne University. While with DSTO, Dr Mansell has worked on the application of information fusion techniques to Naval problems (including mine warfare and combat systems), and he is currently Task Manager for research in artificial intelligence decision aids. His main research interests are information fusion, uncertain reasoning, AI planning, and tactical decision aids.

D.R. Skinner

Maritime Operations Division



David Skinner received a BSc in Physics from the Queen's University of Belfast, Northern Ireland, in 1957, and until 1960 worked with Short Brothers and Harland on the design of aircraft transducers. Since then he has been with MRL Maribyrnong, where he is a Principal Research Scientist and Task Manager for research in artificial intelligence for naval applications. His research experience includes work in theoretical and experimental xerography, laser physics and its applications, remote colorimetry, visual camouflage and sonar surveillance. Research interests include digital image processing, visual perception, and adaptive systems.

K.K. Benke

Maritime Operations Division



Kurt Benke received BSc and MSc degrees in Physics from Melbourne University, a PhD degree in Mathematics/Computing from Deakin University, and a postgraduate Diploma in Applied Statistics from RMIT. At the Kodak Research Laboratory in Melbourne, he carried out theoretical and applied research in physical optics, X-ray physics and electromagnetic scattering. In late 1981 he joined the DSTO, where he completed his PhD studies in 1986 and was involved in research in human and machine vision, mathematical optimisation, and texture analysis. Since 1988 he has been active in sonar data processing, including automatic mine detection, sonar performance assessment, and search pattern optimisation in route surveillance.

Contents

1. INTRODUCTION	1
1.1 Mine Countermeasures Domain.....	1
1.2 AI and MCM Applications.....	2
1.3 Techniques Investigated.....	3
1.4 Report Organisation.....	4
2. LOGIC AND KNOWLEDGE REPRESENTATION	7
2.1 Propositional Logic	7
2.2 Predicate Calculus.....	8
2.3 Expert Systems	9
2.4 Frames.....	10
2.5 Semantic Networks	10
3. REASONING UNDER UNCERTAINTY	11
3.1 Statistical Inference.....	11
3.2 Probabilistic Inference.....	12
3.3 Evidential Reasoning	12
3.4 Fuzzy Logic.....	13
3.5 Probabilistic Logic	14
3.6 Artificial Neural Networks.....	15
4. OTHER TECHNIQUES.....	16
4.1 Updating Decisions.....	16
4.1.1 Nonmonotonic Reasoning	16
4.1.2 Maximum Relative Entropy	16
4.2 Classification.....	17
4.2.1 Cluster Analysis	17
4.2.2 Figure of Merit.....	17
4.2.3 Templating.....	17
5. DISCUSSION AND SUMMARY	18
5.1 General Comments	18
5.2 Specific MCM Applications.....	20
6. REFERENCES.....	21
APPENDIX A: STATISTICAL INFERENCE	24
APPENDIX B: PROBABILISTIC INFERENCE.....	26
APPENDIX C: EVIDENTIAL REASONING	28
APPENDIX D: FUZZY REPRESENTATIONS	33

D.1 Fuzzy Sets	33
D.2 Fuzzy Logic.....	36
APPENDIX E: ARTIFICIAL NEURAL NETWORKS	38
APPENDIX F: GLOSSARY OF TERMS USED IN THIS REPORT.....	41
Fundamental AI	41
Artificial Neural Networks.....	42
Quantitative Inference	43
Decision Updating	47
Classification.....	47

1. Introduction

Artificial Intelligence (AI) is an interdisciplinary science that examines human intelligence by building computational models that emulate what is commonly associated with human intelligence. AI techniques are particularly useful in problems where human performance may be compromised by the volume of information available and the speed of processing which is required. AI is commonly described as being an area of computer science¹ that focuses on the application of symbolic representation and inference to problem solving, as opposed to the more conventional numerical approaches used in traditional computer science programmes.

One concept dealt with in the field of artificial intelligence is how to utilise uncertain and incomplete information. Very little knowledge in the world is precise, certain, or complete. For example the information is incomplete when you know a body of water has been mined, but you do not know how many mines have been laid, where they have been laid, or the type of mines laid. The information is uncertain when you do not know whether, given an opportunity, a mine will detonate, or whether it is defective, or has been rendered inoperative.

Of particular interest to mine countermeasures (MCM) applications is how imprecise information is represented and used for reasoning by a computer. In this report, we summarise an investigation of selected AI techniques which show promise in minewarfare modelling and mine countermeasures. A glossary of terms used is given in Appendix F.

1.1 Mine Countermeasures Domain

Mine countermeasure (MCM) operations comprise four major activities, namely

- ♦ clearance diving: the use of free-swimming Navy personnel to locate, identify and possibly dispose of individual mines and mine-like objects (MLOs),
- ♦ minehunting: the use of specialist craft to locate, identify and possibly dispose of individual mines and MLOs,
- ♦ minesweeping: the use of specialist craft or craft of opportunity (COOPs) to cause mines to explode harmlessly by misleading mine sensors or mine-actuation algorithms, without necessarily locating individual mines, and
- ♦ route survey: the use of specialist craft or COOPs to locate, record the positions of and possibly identify individual mines and MLOs, and to identify safe paths through potential mine fields.

¹ The field of artificial intelligence is considered to spans a diverse range of disciplines, including computer science, mathematics, physics, psychology, engineering, and philosophy.

All MCM activities are subject to considerable uncertainty as to the threat to be countered. Even when the country or organisation responsible for a mine field can be identified, there may be uncertainty as to the type or types of mine laid. In addition, any given modern mine can be configured in many ways, with variations in parameters such as

- ♦ the ship count (a ship count of n means that $n - 1$ ships assessed as targets are allowed to pass unharmed before the mine is 'poised' to explode on the next presumed target),
- ♦ the sensitivities of various detectors and the values of critical time intervals, and
- ♦ the type of algorithm by which the mine-actuation system determines whether it has detected a target.

Finally, the number and location of mines will never be known with any certainty. All of these difficulties are currently handled using probabilistic and statistical methods. In all of the four MCM activities, it is not usual for initial estimates of critical factors to be updated continually on the basis of events that have been experienced, such as the number of mines activated during sweeping. If an unexpected event occurs, however, operations may be stopped whilst revised tactics are considered.

1.2 AI and MCM Applications

In MCM operations, decisions must be made on the basis of previous experience combined with a wide variety of information, some verifiable and quantitative in nature, and some based on tentative assumptions and reports of varying reliability. AI techniques have the potential to supplement existing algorithms under a variety of conditions. These include cases where:

- ♦ An algorithm exists, but with present computing techniques is incapable of running in real time. Here, AI techniques would be used to summarise the conclusions reached by repeated off-line applications of the algorithm. For example, the expectation for the effectiveness of a mine-hunting operation, or for the probability that a mine will operate within the damage radius of a particular target, is currently computed by repeated application of determinate physical models, possibly using some type of stochastic approach. Acquisition of sufficient data for decision-making requires many hours of computation, but, *after this has been done*, the results can be presented in a form, e.g. as an artificial neural network, that can run in a few seconds on a minimal computer.

- ♦ Decisions are based on experience that is difficult to describe quantitatively. In this case, AI techniques can be used to simulate the decision process without the necessity for quantifying, or even describing, the processes involved. An example of this would be the detection of ground mines in sonar displays. An intelligent system can learn, through experience, to make the same decisions that experienced operators have made in a representative selection of operations, and so can present a possibly inexperienced operator with automatic cueing aids. The techniques involved here might be neural networks or fuzzy logic, separately or in combination.
- ♦ Decisions are based on qualitative rules, which may not even have been formulated explicitly, using a wide variety of qualitative and quantitative data and criteria. In such cases, data may be inaccurate, missing, of variable reliability or even contradictory. Here, AI techniques can be used to summarise the data and to present the appropriate commander with estimates of the possible consequences of various options. Such a case might be the selection of the most effective use of assets (clearance diving, minehunters, minesweepers) for clearance operations. This type of problem, based on a (usually complex) set of rules gained from experience, is typical of what are usually called production systems, and additional AI techniques involved here are likely to include fuzzy logic and evidential reasoning.

This report summarises an investigation of techniques considered by the authors to be appropriate to particular minewarfare modelling and mine countermeasures applications.

1.3 Techniques Investigated

Artificial intelligence is not just a single technique; rather, it is a name loosely applied to a large variety of techniques. Often these approaches are intended to represent, to some extent, some of the decisions and assessments made by a human expert in a field of interest. Table 1 shows the techniques (AI and others) that have received some consideration in this report.

These techniques can be divided, for the purpose of the MCM problem domain, into a number of major groups:

- ♦ logical inference - techniques that require full knowledge of the conditions under which decisions must be made, and then perform a (usually complicated) series of operations or calculations,
- ♦ uncertain reasoning - techniques that make allowances for missing, inaccurate and/or inconsistent data in coming to what is intended to be the most probably correct solution,

- ♦ functional relationship - techniques that describe behaviour of interest in terms of contributing factors, usually when the objective relationship is too complicated to be modelled (or perhaps even understood),
- ♦ decision updating - techniques that present a model of the system of interest, and are used to correct this model in the light of new information, and
- ♦ classification - techniques that make a decision on the identity or character of an object or person, using whatever information is available.

Naturally, not all of these techniques are of equal interest to the objectives of this report, and some have been included merely for completeness. As foreshadowed above, the techniques of most interest will be shown to be the evidential-reasoning and fuzzy-logic approaches to uncertain reasoning, and the representation of functional relationships using artificial neural networks. The remainder of the report comprises a brief overview of all the techniques referred to in Table 1, followed by appendices describing in more detail the five most relevant to MCM applications.

Whilst it is not feasible to give a comprehensive overview as to which techniques are suited to particular types of problem, Table 2 gives an indicative set of descriptions and applications.

1.4 Report Organisation

Section 2 of this report outlines some of the standard approaches provided by AI to representing knowledge and inference rules. It is apparent from this section that formal logic and its adaptations are not appropriate to dealing with uncertain and/or incomplete information. However, representation and reasoning with imprecise information usually involves the development of a hybrid system that includes a formal knowledge representation schemes adapted to utilise a specific uncertain reasoning technique.

Section 3 summarises those techniques identified by the authors as being of particular use for representing and reasoning with uncertain and/or incomplete information. These approaches include statistical inference, probabilistic inference, evidential reasoning, fuzzy logic, and artificial neural networks. A detailed examination of these techniques and how they may be applied to MCM problems is found in Appendices A to E.

Section 4 briefly describes other techniques (some borrowed from pattern recognition) that may be useful for dealing with imprecise information, but show less appropriateness to the MCM applications under investigation in this report.


Section 5 is a discussion of the relative appropriateness of the AI techniques investigated to some MCM applications.

Appendices A to E explain in a textbook fashion statistical inference, probabilistic inference, evidential reasoning, fuzzy logic, and artificial neural networks respectively. Included in these sections are examples of how one might apply the techniques to an MCM problem.


Appendix F is a glossary of terms set out in a functional format, describing in an informal manner, the meaning of some of the technical terminology used in AI.

Table 1 - Summary of Techniques Investigated

Logical Inference




Uncertain Reasoning



Functional Relationships



Updating Decisions



Classification




Table 2 - Indicative Applications for Investigated Techniques

AI Method	Description and Indicative Application
Predicate Calculus	A formal logic system applicable when the behavioural rules for a system, and the inputs to the system, are completely known.
Frames	A means of describing numerous examples of related objects (e.g. ships), with known interactions between them.
Semantic Networks	A graphical knowledge representation scheme appropriate when describing objects and complicated relationships between them, such as inheritance of properties, ownership and interactions.
Statistical Inference	A technique for estimating confidence in alternative hypotheses given information on statistical distributions of contributing factors.
Probabilistic Inference	A system of approximate reasoning back from events to causes, given the probabilities of all causes and the probability of the event occurring as a result of each cause.
Evidential Reasoning	A method of determining confidence in alternative hypotheses, given an empirical or subjective assessment of beliefs in propositions that may be incomplete or inconsistent, and may be from different sources and/or expressed in different frames of reference.
Fuzzy Logic	A formal logic system appropriate when information is imprecise and/or when rules for reasoning are approximate.
Probabilistic Logic	A formal logic system that produces estimates of the probabilities of logically provable events, given sets of propositions and events, with empirical or subjective probabilities of their truth.
Neural Networks	A system capable of learning to produce a required set of results for a representative set of inputs, and used to estimate the expected results from different sets of inputs.
Nonmonotonic Reasoning (NMR)	A form of reasoning based on qualitatively ranked statements, using the best available information, and including a process for withdrawing conclusions in the light of new evidence.
Maximum Relative Entropy	A form of logical reasoning based on selecting data so as to minimise the uncertainty of conclusions reached.
Cluster Analysis	A classification technique based on the position in a multi-dimensional parameter space of the properties of a system or event.
Figure of Merit	A classification technique based on algebraic functions of the properties of a system or event.
Templating	A classification technique based on the extent to which the properties of a system or event comply with given criteria.

2. Logic and Knowledge Representation

Artificial intelligence systems are often characterised by their approach to symbolic search and representation. Often this will involve a knowledge base, used to store generic and domain specific information, and an inference mechanism, used to draw conclusions and reason. The inference mechanism searches through the knowledge base looking for solutions or answers to specific problems or questions.

Logical inference requires a complete and precise description of the problem to be solved, and of the conditions that apply for a given attempt at solution. It then uses the conclusions of existing solutions that, when combined using known rules of inference, approach the target solution until the chain of inference leads to what is required. It is thus applicable principally to relatively simple systems with clearly defined rules, such as theorem proving in algebra and geometry. When it is applicable, however, it has the advantage that it is characterised by a result that is known to be valid, consistent and precise.

This section discusses some of the more common forms of knowledge representation, and examines various techniques for producing inference mechanisms.

2.1 Propositional Logic

The term *propositional logic* (Frenzel, 1987) is used to describe what one might refer to as classical logic, and it was therefore one of the first representations schemes used in AI. Here, problems are solved deductively using rules of inference to derive a conclusion, given certain axioms. The form, or syntax, of a statement is rigid and the determination of truth is by syntactic formula manipulation. Propositional logic deals with constant statements (or propositions) known to be either true or false. Legal connectives in the construction of statements are **and**, **or**, **not** and **if**. The overall expressive power of this form of logic is restricted by the simple connectives available. Barr and Feigenbaum (1981) have pointed out that this results in a difficulty in expressing complex concepts.

Propositional logic allows us to express statements like, *if minehunter is in dry-dock, then it is not available for service*. Given the propositions:

X = minehunter is in dry-dock, and
Y = available for service.

The sentence (or properly formed logically expression) can be represented arithmetically as

$$X \Rightarrow \neg Y.$$

When such statements are broken up into combinations of variables and connectives, sentences of propositional logic can be constructed. These sentences can then be manipulated similarly to normal algebraic expressions in mathematics.

2.2 Predicate Calculus

The expressive power of propositional logic is generally insufficient for knowledge representation. *Predicate calculus* (Barr and Feigenbaum, 1981; Frenzel, 1987), an extension of propositional logic, allows one to describe the objects that make up a proposition, and reason about both object and proposition. The expressive power of predicate calculus comes from the way knowledge is represented. Predicate calculus in conjunction with first order logic allows for the association of qualities and attributes with objects, for relationships between sets of objects, and for general statements to be made about objects.

Predicate calculus has a well-defined formal semantics, and its inference rules are sound² and complete³ (Charniak and McDermott, 1985). Like propositional logic, it is a language for representing propositions and rules to generate facts from those given to the system. Predicate calculus consists of predicates that are statements about individuals or objects, their properties, and their relationships with other objects, which return a true or false value. Predicate calculus also allows the manipulation of quantified statements such as *all current mines have acoustic wake-up*. This may be expressed in predicate calculus using the quantifier \forall , meaning *for all*, and the variable X , as

$$\forall X, \text{CurrentMines}(X) \Rightarrow \text{AcousticWakeUp}(X).$$

Similarly, the expression, *there is an FFG-class vessel that is friendly*, may be expressed using the quantifier \exists , meaning *there exists*, and the variable X , as

$$\exists X, \text{FFGClass}(X) \wedge \text{Friendly}(X).$$

Reasonably complex expressions and assertions can be made when presented using formal expressions in sentences of first-order logic.

The use of predicate calculus as a knowledge representation scheme in AI has met with mixed results. Although resolution will always provide a correct answer if all information is correct and an answer exists, the system is very general and clumsy. When the problem becomes non-trivial, there is a combinational explosion in the

²Describing an inference for which, given a set of propositions and an inference rule, every inference follows the inference rule (Mercadal, 1990).

³Being able to derive all possible inferences from a set of propositions (Mercadal, 1990).

number of alternatives to be investigated. In an effort to constrain the search in large databases, heuristics have been employed to choose which approach would be most feasible. Another major drawback in the use of first-order predicate calculus is the restriction placed on the knowledge representation scheme by not allowing relationships between *predicates* (i.e. assertions), beliefs, temporal relations or statements of possibilities. Predicate calculus is a convenient representation for facts and rules of inference, provided the domain can be adequately captured by the knowledge engineer (or person who interacts with a domain expert in order to acquire relevant facts and relationships among facts to be built into an AI system).

2.3 Expert Systems

The logic representations discussed so far have consisted of a finite set of formally defined formulae and statements. This has proved restrictive for application to real-world applications, where constraints can be ill-defined or non-existent. This deficiency resulted in the development of a variety of schemes generally known as *expert systems* (also known as expert systems or knowledge based systems) (Barr and Feigenbaum, 1981; Tanimoto, 1987), which are computational models used for implementing search algorithms and for modelling human problem solving. A typical system would consist of a set of production rules, a working memory, and a control cycle. The production rules are cast as a group of condition-action pairs of the form "If this *condition* holds, then this *action* is appropriate." Their actions are specifically designed to alter the contents of the working memory, which holds a *world model* (description of the problem) in a buffer-like data structure. The control structure of an expert system operates on a subset of the working memory for conflict resolution, identifying conflicts between the real and current worlds, and effectively selecting the production rules to be executed one at a time.

For example, an expert system may hold in its working memory a representation of the environment, part of which includes the statements {*Sonar(active)*, *Target(nil)*}. This part of the world description may be manipulated by a number of production rules such as *if sonar contact, then target located*.

Expert systems are most often used in AI programs to represent a body of knowledge about how people do a specific task. The inherent disadvantages of expert systems is that their strong modularity and uniformity result in a high inefficiency in problem solving. Although situation-action knowledge can be expressed naturally this way, algorithmic knowledge cannot, making the control logic difficult to follow. Also, the application-inspired design tends to make such a system very problem-specific. Three types of implementation of expert systems are described below.

2.4 Frames

Frames (or *schema*) are used to group information about particular objects and situations. A frame can be viewed as a static data structure used to represent well-understood, stereotyped situations, with the inter-relationships of objects represented as slots in a frame, and values of the properties stored in the slots. An interesting feature of a frame is its ability to determine whether it is applicable to a given situation and, if not, to transfer control to a more appropriate frame. Each individual frame can be viewed as a data structure, similar in many respects to the traditional record, that contains stereotyped entities.

For example, in a frame-like language a submarine may look like this:

Generic SUBMARINE Frame

Description: Vessel, Boat.
 Class: Delta, Collins.
 Alliance: Friend, Foe, Neutral.
 Type: SSBN, SSK.

Sonar-Contact Frame

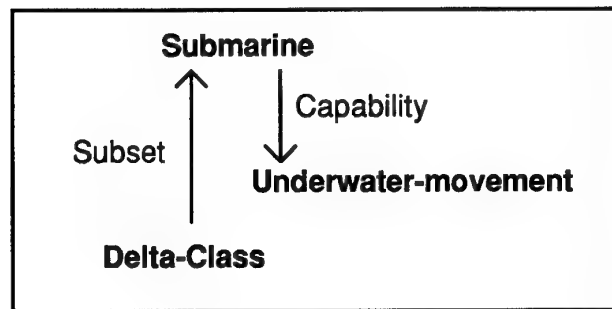
Description: Vessel.
 Class: Delta.
 Alliance: Foe.
 Type: SSBN.

Although research into frames is continuing to find new applications, it is unlikely that they will have much application to MCM activities, since such well-defined problems are generally already treated by proven algorithms.

2.5 Semantic Networks

The *semantic network* takes a set of logical predicates and represents them graphically, with nodes corresponding to facts or concepts, and arcs (or links) in the graph instead of predicates to indicate relationships. An algorithm for reasoning about associations within the domain then simply needs to follow the links. In addition, semantic networks implement inheritance, i.e. certain links in the network indicate class membership and allow properties attached to a class to be inherited by all members of the class.

For example, the following simple semantic net represents the statements *a Delta-class vessel is a submarine*, and *a submarine is capable of underwater travel*.



3. Reasoning Under Uncertainty

One feature that all of the schemes for logical inference have in common is the need for a complete and accurate world picture. Such systems apply universally valid rules to absolutely certain facts to deduce more facts of absolute certainty. This requires a model of all objects and the rules governing every possible interaction between them.

In the real world, complete information about the environment is generally unavailable. One must therefore take into account the varying degrees of uncertainty inherent in any particular environment, and make the best possible decision with the evidence available.

There are many different functions that a method of reasoning under uncertainty may serve, for example:

- ♦ representation of degree of belief,
- ♦ evaluation of the strength of an argument,
- ♦ application of rules of general but not universal validity,
- ♦ inference based on uncertain, incomplete, or qualitative concepts.

The consideration of such issues has led to the development of a number of schemes for *uncertain reasoning*, and some of the more significant are discussed in the following section.

3.1 Statistical Inference

Statistical inference (Davis, 1990; Flachs, Jordan and Carlson, 1988) is one of the simplest forms of uncertainty reasoning, being based on the primary statistical concept that a population of events may be adequately represented by a sub-set of itself. In making this assumption, it is important to be able to ascertain whether information

available represents a significant event to within a specific confidence level. Statistical reasoning is well suited to this, and hence is used in radar tracking and pattern recognition or classification, such as the example in Appendix A.

This technique is clearly useful when precise probabilities for the test and null hypotheses (in this case, detection and false-alarm probabilities) are known (as shown by its application to the fusion of sensor data in Appendix A). It does not appear to be usefully applicable to less structured world models.

3.2 Probabilistic Inference

Although predicate calculus is a widely accepted form of knowledge representation and includes inference procedures representing a form of logical deduction, in practice human reasoning uses terms such as *probably*, *usually* and *occasionally*, etc. demonstrating that its patterns are intrinsically probabilistic. This is not to say that the underlying logic of such patterns cannot be axiomised. In fact, probability can be viewed as a generalisation of predicate calculus, where the truth value of a proposition, given some evidence, is no longer a Boolean value, 0 (false) or 1 (true), but is generalised to be the real interval between 0 and 1, probability *in this context* being a measure of belief in a proposition. In *probabilistic inference* (see Appendix B) all relevant inference paths that connect evidence to hypotheses of interest must be examined and combined, in contrast with predicate calculus, where it is sufficient to establish a single path between the axioms and the theorems of interest.

In many real-world scenarios, uncertainty methods that are Bayesian based will require the use of some probabilities that are not available and must be estimated. If the sample space of the probabilities is well understood, then the estimation theory can be matched to give estimates close to the true probabilities. However, one must always realise that these are only estimates and hence will contribute to the degree of uncertainty in the final decision made. It follows that, Bayesian statistics are most useful in drawing conclusions from the behaviour of a system where the probabilities of events are well understood, for example, robot control and sensor fusion.

3.3 Evidential Reasoning

In a system that is uncertain or ill-defined, a single body of evidence (BOE) may give the degree to which any one proposition should be believed. However the precise degree of belief that should be accorded every environmental proposition cannot be calculated. The amount of ignorance a BOE contains is hence an important component in the reasoning process. It is for this reason that Bayesian point probabilities are often an inadequate form of reasoning from evidence. The requirement that each probability be assigned a precise value in Bayesian theory leads to confusion as to whether a low probability implies that there is no particular reason to believe that this BOE is true, or that there is good reason to believe it is false. In Bayesian theory, this

form of uncertainty can only be represented by a second order probability (i.e. the probability that the first order probability value is true).

This confusion can be avoided by the implementation of *evidential reasoning*, a set of techniques based on *Dempster-Shafer* (D-S) theory, which is a mathematical theory of evidence conceived by Dempster (1968) and further modified by Shafer (1976). Being a departure from classical probability theory, D-S reasoning uses information that is typically uncertain, incomplete and error-prone. D-S theory maintains the association between the measure of belief and disjunctions of events rather than forcing probabilities to be distributed across a set of possibilities. The result is that one need no longer assume that all data are available and being utilised. Dempster-Shafer theory is a way of capturing both the first and the second order information using only first order numbers.

Shafer [1976] writes, "the additive degrees of belief of the Bayesian theory correspond to an intuitive picture in which one's total belief is susceptible to division into various portions, and that intuitive picture has two fundamental features. First, to have a degree of belief in a proposition is to commit a portion of one's belief to it. And secondly, whenever one commits only a portion of one's belief to a proposition, one must commit the remainder to its negation. The obvious way to obtain more a flexible and realistic picture is to discard the second of these features while retaining the first."

Because evidential reasoning is considered highly relevant to MCM decision making, it will not be described in this brief summary, but is the subject of Appendix C.

3.4 Fuzzy Logic

Another approach to handling imprecision in decision making is through the concept of *fuzzy sets*, with their extension to *fuzzy logic* (Zadeh, 1965, 1983). (This latter term is unfortunate, but too well established to be changed - fuzzy logic is not fuddled thinking, but clear thinking about imprecise concepts.) Fuzzy sets were developed in order to handle *linguistic variables* (or predicates); for example, an observer might refer to the weather conditions as "fairly windy", and this may be all the information available. It would clearly be unreasonable to assign either a precise value of wind speed, in metres per second, to such a variable, or to assign a given range of speeds. Further, it may be necessary to build into a logic system condition-action clauses like "if the weather is *very windy* then sonar detection becomes *quite inefficient*". It would be an undesirable and unconvincing algorithm that estimated mine-detection probability as, say, 50% at wind speeds up to 10 m/s, and as 10% for all other speeds (including, e.g. 10.01 m/s), which would be the case if one assigned a precise boundary to the speeds corresponding to *windy*, and a precise value to the descriptor *efficient*.

In traditional, or *crisp*, set theory, an object either belongs in a set, or it doesn't. Thus, weather with a given wind speed is either a member of the set *windy*, in which case it has a membership value of 1, or it is not a member, and has a membership value of 0.

In fuzzy set theory and fuzzy logic, this membership value is replaced by a *membership function*, in the real interval 0 to 1, which describes the extent to which an object belongs to the set. (The membership function is not the same as the subjective probability of Bayesian inference or the degree of belief of Dempster-Shafer theory, although, in some ways, it resembles both.) Thus, an object may have non-zero membership functions in both a set and its complement, e.g. *windy* and *not-windy* (or *calm*). It then becomes possible to handle compound linguistic concepts, such as "not windy and not calm", which have meaning in normal thinking but not in crisp set theory. Extension of this concept to the set-theory representation of condition-action clauses effectively eliminates discontinuities of the type referred to in the previous paragraph.

Dempster-Shafer theory and fuzzy logic have a degree of similarity, in that they are both used to represent uncertain and conflicting information. However, Dempster-Shafer theory deals primarily with the combination of information from different (and possibly conflicting) sources, whilst fuzzy logic deals with imprecise measurements and qualitative concepts. As with Dempster-Shafer theory, fuzzy logic will be considered later, in Appendix D.

3.5 Probabilistic Logic

Probabilistic logic, developed by Nilsson (1986) is a semantic generalisation of ordinary first-order logic. Each proposition of interest is given a truth value representing the probability that it is true, and a set of possible worlds is established (i.e. if there were one proposition, then there would exist two possible worlds, one where the proposition is true and the other where the proposition is false). These propositions can be true in some worlds and false in others, as long as they are in different combinations, and each possible world must contain a unique and consistent set of propositions. This would imply that, if there were L propositions, then the number K of possible worlds could be as high as 2^L . However, there are typically fewer than this, as some combinations of true and false propositions are inconsistent.

Nilsson uses a matrix notation for the representation of probabilistic logic. In a simple situation with point probabilities, the relationship between the L -dimensional column vector Π representing the probabilities of the propositions, the K -dimensional column vector P representing the probabilities of the various worlds, and the $L \times K$ matrix V representing the truth values for the propositions in these worlds is simply

$$\Pi = V.P$$

Nilsson extends this concept to reasoning with uncertain beliefs, when the probabilities of the possible worlds are not usually given, and one must determine them from the available information. Using a base set of beliefs with associated probabilities, we can deduce a new proposition and its associated probabilities. We

now know V and Π and can solve the matrix equation for P . Nilsson terms this operation probabilistic entailment and uses it to calculate the probability of a proposition being true or false and the probability of an operator being in a given possible world. However, this technique does not appear to be strongly relevant to MCM operations, since it deals with beliefs and probabilities rather than condition-action systems.

3.6 Artificial Neural Networks

In cases where it is not possible, for reasons of complexity or lack of knowledge, to describe the behaviour of a system of interest as an explicit function of the contributing factors, it is frequently convenient to use a simple approximation to the relationship. In the past, because of the sheer weight of computation involved in any non-linear least-squared-error type of calculation, the preferred method has been *multiple linear regression*. With the relatively recent advent of *artificial neural networks*, however, it is now possible to describe non-linear systems in a convenient way. This is not, strictly speaking, an AI technique, since it embodies only information on effects, rather than the mechanisms that cause them, but it is so useful a component of intelligent systems that it is described in more detail later, in Appendix E.

Examination of data from recent trials in Jervis Bay (Neill, 1991) reveals that the measured navigational accuracy of a sonar platform correlates with system and environmental variables (such as ship speed, wind speed, wind direction, sonar orientation etc). A predictive model linking system and environmental data to navigational accuracy could conceivably be used to flag unfavourable operating conditions, leading to possible postponement of the mission or appropriate operational changes (resulting in a saving in operating costs).

The results of a recent study, which investigated neural networks as a tool in mathematical modelling, suggests that the hover radius of the MHI minehunter could be modelled as a function of system and environmental variables (Benke 1993). A correlation coefficient of 99% was obtained between predictions and measurements when applied to new data, as opposed to 56% by multiple linear regression. The approach was shown to be effective for modelling the quantitative effect on performance of different human operators. Other applications include the identification and classification of ship and mine signatures, and as an integral part of a mine logic system.

4. Other Techniques

The techniques outlined in this section are taken from AI and pattern recognition. They are grouped under a catch-all heading as they are less likely to be applied to MCM problems (as described in section 1.1).

4.1 Updating Decisions

The following is a brief description of two mechanisms that are used for revising a model of a universe of discourse, rather than setting up a new model. These are considered to be of marginal interest in the early development of an MCM application of artificial intelligence.

4.1.1 Nonmonotonic Reasoning

The most compelling reason for using first-order logic as a framework for representing and combining information is that logical inferences based on unambiguously true statements never result in invalid conclusions. The most significant disadvantage is that it cannot effectively accommodate uncertain, incomplete or inconsistent information. A *nonmonotonic reasoning* (NMR) system (McDermott and Doyle, 1980) handles uncertainty by making, at each decision point, what is believed to be the most reasonable assumption in light of the available evidence. If, at a later time, an assumption is found to be erroneous, because of either new evidence or the discovery that the assumption led to an impossible conclusion, the system changes the assumption and all the conclusions that rely on it. Thus, in contrast with first order logic, the number of possible statements from a set of assumptions does not necessarily grow monotonically with the addition of new information. Since information can be retracted in NMR systems in the light of new information, it is important to keep track of all deduced knowledge; when an assumed fact is withdrawn, all conclusions dependent on it must be re-examined and possibly withdrawn.

4.1.2 Maximum Relative Entropy

Maximum relative entropy inferencing (also known as cross-entropy inferencing or minimum-information updating) is a method for updating a probability distribution in the light of new information on currently defined propositions (Jaynes, 1982). Maximum entropy is a dynamic theory where previously mentioned theories are static; a dynamic theory is concerned with how a belief should change in the light of new information, while a static theory is concerned with consistent conditions for degrees of belief at a given time.

4.2 Classification

The following comprises a brief description of three methods of classifying objects or events on the basis of uncertain information. They are included mainly for completeness, and it is not considered that they will have significant application to foreseeable MCM problems.

4.2.1 Cluster Analysis

Cluster analysis (Everitt, 1977) is generally used for classification analysis based on multi-parameter similarity. This is achieved by sorting observations into natural groups based on the estimates of pair-wise and cluster-wise similarities. Observations are cast into a non-dimensional form, and assembled into a multi-parameter space, where one of a variety of techniques is used to create a *resemblance matrix* defining the similarity between each pair of objects (or events).

4.2.2 Figure of Merit

Figure of merit calculations are similar in principal to cluster analysis, but differ in detailed application. They are used by LOCE (Limited Operations Capability Europe) (Llinas, 1989) to fuse information from electronic intelligence reports, photo-interpretation, target data messages, and free text. This involves calculating the degree of similarity between two entities using their attribute vectors. The LOCE system uses a self-correlation process involving location, frequency, pulse width, pulse repetition interval and time, followed by a cross-correlation process to associate new data with higher level entries.

4.2.3 Templating

Templating is often used in decision fusion by first establishing preset logical or numerical criteria to determine if a certain set of observations supports an event or conclusion. One or more parametric or nonparametric observations is collected, possibly over a period of time, and, using weighted thresholding, Boolean templates, and hierarchical event profiling, a declaration is made of whether an event or object matches an expectation.

5. Discussion and Summary

5.1 General Comments

Examination of the naval literature (see, for example, Pollaers 1985, Hartman 1988) indicates that the principal aims of artificial intelligence in maritime operations include:

- ♦ the selection of weapons options to produce maximum effect on target,
- ♦ the selection of tactics to produce maximum effect on the battlefield,
- ♦ the facilitation of naval warfare mission planning,
- ♦ the alleviation of operational problems due to manpower shortages and frequent staff re-assignments,
- ♦ the enhancement of multi-sensor integration to reduce information overload, and
- ♦ the improvement of reaction times against missile threats.

A number of techniques for dealing with uncertain and incomplete information have been investigated in this report. Many of these techniques are treated in the AI literature, and involve statistical and probabilistic inference, evidential reasoning, fuzzy logic, artificial neural networks, and nonmonotonic reasoning. Some approaches, however, such as maximum relative entropy, cluster analysis, and figure of merit, originate from the pattern recognition literature. All of these approaches are well suited to specific problem types. Hence, this report does not evaluate these approaches, rather it investigates their appropriateness to specific MCM applications.

The investigation has highlighted evidential reasoning, fuzzy logic, and to a lesser extent, artificial neural networks, as they were deemed by the authors to demonstrate facilities most useful to particular MCM problems. Evidential reasoning can be used for representing uncertain and incomplete information, and provides a powerful range of operation for manipulating bodies of evidence. An appropriate application of evidential reasoning in the MCM domain would be involve the filtering and processing of the large quantities of information available to an MCM commander. Evidential reasoning can be used to combine bodies of evidence, emphasising common attributes, and de-emphasising contradictory information. It can then provide the MCM commander with a detailed or summarised report on the information (depending on the individuals requirements). It is envisaged that evidential reasoning would best serve the MCM domain as an aid to operator by arranging information in this way, not replacing the operator in the decision process.

Fuzzy logic deals with a different kind of uncertainty from that appropriate to evidential reasoning - the uncertainty here is mainly in the quantitative values of

factors, rather than in confidence in the existence of specified conditions. Fuzzy logic can also handle relationships represented in terms of vague concepts. It is appropriate for presenting information derived from fuzzy or 'crisp' algorithms where the input data are inherently inexact, and it can therefore be used for the development of tactical decision aids. As an example, the decision as to the detailed tactics (or 'stages') to be employed by clearance divers in searching for mines is made using a crisp algorithm that has among its inputs a number of very approximate physical measures and some rather arbitrary thresholds. The paradigms used by fuzzy logic are ideally suited to following through the algorithm and presenting to a decision maker confidence in the applicability of each tactic. The related field of fuzzy control systems is a mature technology, suited to emulating the behaviour of experienced operators in activities as diverse as focusing cameras and steering power boats. This should have many applications in MCM operations, particularly where it is desirable to replace an operator by an automatic controller under hazardous conditions.

Artificial neural networks are best suited to classification problems in domains where good training data are available. The parabolic-exponential model fitted to the lateral range function of a sidescan sonar by regression analysis is a sufficient approximation under some operational conditions. There may be cases, however, where consideration could also be given to a completely distribution-free method, such as that offered by the use of an artificial neural network. The advantage in this case is the fact that no *a priori* model is assumed for curve fitting and the approach is therefore more generalised. The development of an autonomous cueing aid for sidescan sonar during route surveillance can also be enhanced by using a target classifier (neural network) to process the data from the outputs of tuned spatial filters.

The solution of MCM problems, or the provision of advice on the likely effectiveness of possible courses of action, is typical of the sort of application for which expert systems, are well suited. Expert systems in their traditional forms have sometimes experienced difficulty in expressing explicitly the expression of algorithms, and also accounting for uncertain, missing or contradictory data. Advances in the incorporation of fuzzy logic into expert systems have led to considerable improvement in the handling of approximate or linguistic data, and the more recent application of artificial neural networks to expert systems has allowed for the implementation of algorithmic knowledge in a real-time manner. Problems with missing or contradictory data do not yet appear to have been solved within expert systems, although evidential reasoning has existed for some time as an appropriate tool for handling such data.

It appears that the most likely form for an AI system to take, for the solution of MCM problems, will be one including fuzzy algorithms and neural networks, and that the incorporation of evidential reasoning in such a system would be worth investigation.

5.2 Specific MCM Applications

As discussed above, the clearance diver operates on estimates of environmental conditions such as sea-bed type and underwater visibility. Such information is used to predict the consequences of each of several possible tactics in terms such as sea-bed coverage rates for a given clearance level, and a choice of method is made according to criteria that depend on the circumstances, e.g. minimum risk to the divers or maximum clearance level. In short, a predictive model of the performance of a clearance diver would contain little in the way of algorithmic analysis, but a significant amount of qualitative rule-based decision making.

The operation of minehunting is a combination of predictable manoeuvring and stochastic processes. The MCM vessel will usually adhere to a planned route, but may diverge from this to deal with any mine or MLO that is detected, then return to its original course. Since the disposition of mines in a hostile minefield is not known, it is difficult to analyse likely behaviour other than by a probabilistic model, such as a Monte Carlo simulation, using multiple runs to achieve reasonable estimates of probabilities. There are, however, algorithms, tables and nomograms available for the planning of operations.

Minesweeping and route survey resemble each other in that they employ similar procedures, and indeed sometimes make use of the same vessels. These operations are eminently predictable, and both consist of one or more scans in a regular pattern over a pre-selected area. Under normal conditions, they can be planned in advance using well understood computational aids.

MCM vessels will often be operating under difficult conditions, including adverse weather and hostile activity. Such conditions are likely to result in vital data inputs being lost. For example, a route-survey vessel may rely on short-range sonar to determine the relative position of a sidescan sonar tow-fish. If the short-range sonar becomes inoperative due to accidental damage or hostile activity, the position of the tow-fish may need to be input on the basis of an inaccurate technique, or even a rough estimate. Under conditions where a great deal of the necessary information is lost, the vessel commander may have to resort simply to using a best estimate, based on previous experience, for deciding on detailed tactics. The algorithm for deciding tactics may thus be reduced to what is effectively a qualitative rule-based decision.

It can be seen from the above brief summary that MCM operations always occur under conditions of some uncertainty, and, at least in the case of clearance diving, the choice of tactics may be made using data that can never be better than approximate. In bad weather or combat conditions, uncertainties for all operations may be compounded by the loss of normally accurate quantitative information, such as navigational data. Whilst statistical and probabilistic methods offer a means of overcoming some of these difficulties, AI techniques would seem to have considerable potential for assisting decision making by using all of the available information, no matter how sparse, inaccurate or inconsistent. Further, AI techniques could possibly

be used for the continuous revision of assumptions responsible for tactical decisions, and for evaluating possible changes in tactics.

6. References

A. Barr and E.A Feigenbaum (1981), *The Handbook of Artificial Intelligence Volume 1*, Stanford, California: Heuristech Press.

K.K. Benke (1993), *Modelling Hover Performance of Inshore Minehunters by Neural Networks*, DSTO Materials Research Laboratory, Technical Report No. MRL-TR-93-30. RESTRICTED.

Z. Chair and P.K. Varshney (1986), "Optimal Data Fusion in Multiple Sensor Detection Systems", *IEEE Trans. Aerospace Electron. Syst.*, vol. AES-22, pp. 98-101.

E. Charniak and D. McDermott (1985), *Introduction to Artificial Intelligence*, Reading: Addison-Wesley.

M. Cohen, D. Schum, A. Freeling and J. Chinnis (1985), "On the Art and Science of Hedging a Conclusion: Alternative Theories of Uncertainty in Intelligence Analysis", Tech. Rep. 84-6, Decision Science Consortium, Inc. Virginia.

E. Davis (1990), *Representations of Commonsense Knowledge*, California: Morgan Kaufmann Publishers.

A. Dempster (1968), "A Generalization of Bayesian Inference", *J. R. Statist. Soc. B*, pp. 205-247.

B. Everitt (1977), *Cluster Analysis*, London: Heinemann Education Books.

G.M. Flachs, J.B. Jordan and J.J. Carlson (1988), "Information Fusion Methodology", *Proc. SPIE Conf. on Sensor Fusion*, vol. 931, pp. 56-63.

L.E. Frenzel (1987), *Crash Course in Artificial Intelligence and Expert Systems*, Indianapolis, Howard W. Sams & Co.

T.D. Garvey (1987), "A Survey of AI Approaches to the Integration of Information", *Proc. SPIE Conf. Infrared sensors and Sensor Fusion*, vol. 782, pp. 68-82.

M.L. Ginsberg (1993), *Essentials of Artificial Intelligence*. California: Morgan Kaufmann.

P.J. Hartman (1988), "Practical Applications of Artificial Intelligence in Naval Engineering", *Naval Engineers J.*, pp 32-40.

E.T. Jaynes (1982), "On the Rationale of Maximum-Entropy Methods", *Proc. IEEE*, vol 76, pp 939-952.

A. Kandel and M. Schneider (1989), "Fuzzy Sets and their Applications to Artificial Intelligence". In *Advances in Computers*, Vol: 28, edited by M.C. Yovits, Academic Press, Inc., California.

J. Llinas (1989), "Advanced Techniques for Data Fusion and Multi-Sensor Correlation", Seminar notes, Sponsored by Technology Training Corporation, Melbourne, 24-25 July.

D. McDermott and J. Doyle (1980.), "Non-Monotonic Logic I", *Artificial Intelligence*, vol 13, pp. 41- 72.

D. Mercadal (1990), *Dictionary of Artificial Intelligence*, New York: Van Nostrand Reinhold.

B. Muller and J. Reinhardt (1990), *Neural Networks*, Berlin: Springer-Verlag.

R.A. Neill (1991), Sea trials of sidescan sonars mounted on AM Brolga and AM Salvatore V, Jervis Bay, NSW, December (Unpublished).

N.J. Nilsson (1986), "Probabilistic Logic", *Artificial Intelligence*, vol. 28, pp. 71-87.

J. Pearl (1988), *Probabilistic Reasoning In Intelligent Systems*, California: Morgan Kaufmann Publishers, Inc.

J. Pollaers (1985), "Expert Systems", *J. Australian Naval Institute*, vol. 11, pp. 41-47.

D.E. Rumelhart and J.L. McClelland, Eds. (1987), *Parallel Distributed Processing*, Cambridge, Massachusetts: MIT Press.

G. Shafer (1976), *A Mathematical Theory of Evidence*, Princeton, NJ: Princeton University Press.

G. Shafer (1981), "Constructive Probability", *Synthese*, vol. 48, pp. 1-60.

P.K. Simpson (1990) *Artificial Neural Systems*, New York: Pergamon.

S.L. Tanimoto (1987), *The Elements of Artificial Intelligence: An Introduction using LISP*, Maryland: Computer Science Press.

S.C.A. Thomopoulos, R. Viswanathan and D.C. Bougoulas (1987), "Optimal Decision Fusion in Multiple Sensor Systems", *IEEE Trans. Aerospace Electron. Syst.*, Vol. AES-23, No. 5, pp. 644-653, September.

L.A. Zadeh (1965), "Fuzzy Sets", *Information and Control*, vol. 8, pp. 338-353, .

L.A. Zadeh (1973), "Outline of a New Approach to the Analysis of Complex Systems and Decision Processes", *IEE Trans. Syst. Man Cybern*, vol. SMC-3, pp. 28-44.

L.A. Zadeh (1983) "The Role of Fuzzy Sets in the Management of Uncertainty in Expert Systems", *Fuzzy Sets and Systems*, vol. 11, pp. 199-227.

Appendix A: Statistical Inference

As an example of statistical inference, consider the case where a mine detects active sonar transmission from a passing vessel, and it is necessary to determine whether the sonar transmitter is of Type A or Type B, which are known to differ in pulse repetition interval (PRI). However, the probability distributions, $p_A(i)$ and $p_B(i)$, of the sonars operating on particular nominal PRIs overlap as shown in Fig. A.1. Given that we observed a PRI of i_o , we wish to compare the two hypotheses H_0 (the sonar is of Type A) and H_1 (the sonar is of Type B).




Figure A.1: The overlapping pulse repetition interval (PRI) probability distributions for sonars of types A and B.

The strategy is to select a critical value, i_c , and to make the assumptions

$$\begin{aligned} i_o \leq i_c &\Rightarrow H_0 \text{ (the sonar is of type A)} \\ i_o > i_c &\Rightarrow H_1 \text{ (the sonar is of type B)} \end{aligned}$$

Hence the probabilities of incorrect identification are:

$$\begin{aligned} \alpha &= P(i_o \leq i_c | H_1), \quad \text{the probability of selecting } H_0 \text{ given that the true situation is } H_1 \\ \beta &= P(i_o > i_c | H_0), \quad \text{the probability of selecting } H_1 \text{ given that the true situation is } H_0, \end{aligned}$$

where α and β are the shaded areas shown in the figure.

From this, it is possible to show that, if the expected number of occurrences of types A and B are n_A and n_B respectively, and that the costs of single failures to identify the types are C_A and C_B , then the expectation for the total cost, C_T , of all failures is given by:

$$C_T = \beta n_A C_A + \alpha n_B C_B \quad (\text{A.1})$$

and that this will have a minimum value when i_c satisfies the equation:

$$n_A C_A p_A(i_c) = n_B C_B p_B(i_c) \quad (\text{A.2})$$

The solution of this equation for Gaussian probability-density functions is trivial.

In another application, Chair and Varshney (1986) have considered a problem in data fusion (the combination of data from different sources), where sensors make decisions independently of each other before sending their results to a central fusion module for correlation. An optimal fusion rule is derived for the likelihood ratio (LR) test, the ratio of the probability of some pool of evidence being true, given a certain hypothesis, to the probability of it being true given the negation of the hypothesis. This turns out to be a weighted average of the various sensor decisions, where the weights are derived from the individual sensor false alarm and detection probabilities. This approach requires exact knowledge of the *a priori* probabilities of the test hypothesis, or the assumption that all null hypotheses are equally likely.

Thomopoulos, Viswanathan and Bougoulas (1987) derive an optimal decision scheme that has each sensor making an independent decision based on an LR test, and the fusion centre making a further LR test during correlation of the decisions. This information fusion algorithm is applied to two systems, the first where various sensors transmit their decisions only, and the second where they transmit both the decision and the degree of confidence with which it was made. If all sensors are operating under the same conditions, this test will have a higher detection probability than that of the individual sensors; however, in the case of disparate sensors, the system performance is dependent on how different the operational conditions of the sensors are from each other.

Appendix B: Probabilistic Inference

Bayes' Rule of conditioning (Tanimoto, 1987) is the fundamental means of calculating the probability of a hypothesis using measured supporting evidence. Although it is formally defined using *a priori* probabilities, it is often used to upgrade beliefs in a hypothesis based on new evidence. Bayes' Rule simply states that if there is an exclusive and exhaustive set of hypotheses (causes) for an event that has occurred, then the probability that a particular cause was responsible is proportional to the product of the probability of that hypothesis being true and the probability of the event occurring under that hypothesis, that is

$$P\{B_i | A\} = \frac{P\{A | B_i\} P\{B_i\}}{\sum_j P\{A | B_j\} P\{B_j\}} \quad (\text{B.1})$$

where $P\{X\}$ is used for the overall probability of X occurring, and $P\{X | Y\}$ is used for the probability of X occurring given that Y has occurred.

In practice, Bayesian theory is applied by first selecting one event whose outcome is precisely known, and then, using the rules of the theory, calculating the desired probabilities. For example, if A is an observable event, and $\{B_1, B_2, \dots, B_k\}$ is a set of mutually exclusive, exhaustive hypotheses, then one could calculate the probabilities $P\{B_i\}$ and the likelihood $P\{A | B_i\}$ for all i . One can then use Equ. B.1 to calculate the predictive probability $P\{A\}$. Similarly, if sufficient information is available to calculate any two of these probabilities, Bayes' rule can be used to find the third.

As an example, in the detection of a mine-like object through route survey, the performance of the sensor (i.e. the combination of the sonar and any associated target-detection algorithms), can be described by a probability matrix of the type shown in the figure below, where $P\{D_i | O_j\}$ represents the probability that a declaration of (interpretation as) an object of type i will be made, given that the actual object is of type j . This might be, for example, $P\{\text{rock} | \text{Mk-84 mine}\}$.

		Actual Object Type			
		O_1	O_2	L	O_m
Declaration of Type made by Sensor	D_1	$P\{D_1 O_1\}$	L	L	L
	D_2	M	O		
	M	M		O	
	D_n	M			$P\{D_n O_m\}$

In addition to the probability matrix, the Minewarfare Pilot Officer, or other interpreter, will have *a priori* information, from previous surveys, intelligence reports etc., on the actual probabilities, $P\{O_j\}$, of the occurrence of given object types. The

Bayesian equation, Eqn. B.1, can then be used to combine the probability matrix and the *a priori* probabilities to give *a posteriori* probabilities such as $P(O_i | D_j)$, the probability that the object is of type i , given that the sensor has declared it to be of type j .

Similarly, information provided by a number of sensors can also be fused using a multi-variable form of Bayes' equation (Pearl, 1988).

When the necessary probability values for computation are not known, the principle of insufficient reason can be applied. This is simply explained by Garvey (1987): "If the probability of a disjunction of events is known, but the probabilities of the individual components are not, and there is no particular reason to expect that one event is more likely than any other, then the principle of insufficient reason dictates that equal probabilities, totalling to the original probabilities, be assigned to the individual components". Alternatively, a more sophisticated approach available is the maximum entropy principle, which selects probability values by maximising the entropy (or the degree of disorder) of the assignment. This corresponds to making a *minimal commitment* to the estimation of unknown probabilities (Pearl, 1988).

The usual formulation of Bayes' rule, given by Equ. B.1, can be used to obtain the odds-likelihood ratio formulation, by dividing the rule for one hypothesis by the rule for a second hypothesis. This form is often useful when subjective probability judgments (being made by a human) are required, as it is often more intuitive for an assessment of a likelihood ratio to be made, than for a straight conditional probability judgment (Cohen, Schum, Freeling and Chinnis, 1985).

Appendix C: Evidential Reasoning

Evidential reasoning in general, and Dempster-Shafer theory in particular, is used to assess the effect of all pieces of available evidence on a hypothesis, making use of domain-specific knowledge. A propositional space called the *frame of discernment* is used to define a set of basic statements, exactly one of which may be true at any one time, and a subset of these statements is defined as a *propositional statement*. For example, in the case of an intelligent ground mine, the frame of discernment, θ_A , might represent every type of vessel that could influence it, i.e.

$$\theta_A = \{a_1, a_2, \dots, a_n\} \quad (C.1)$$

where one of the basic statements a_i might be "the vessel is a Delta class submarine". A propositional statement A_i might be "the vessel is a submarine", that is, the proposition is the subset of θ_A containing all a_j that nominate different classes of submarine.

In much the same way that one may, given sufficient information, compute probabilities (summing to unity) for all possible combinations of situations of interest, one may assign values (again summing to unity) to one's beliefs in all possible propositional statements in a frame of discernment; these values, $m_A(A_i)$, are known as *masses*, and the process is called a *mass distribution*. This may be written as


$$\sum_{A_i \subseteq \theta_A} m_A(A_i) = 1 \quad (C.2)$$

where the domain of A_i is the set of all possible subsets of θ_A , i.e. the power set 2^{θ_A} . Any proposition assigned a non-zero mass is called a *focal element*, and the mass assigned to the empty set ϕ is zero, since, by definition, at least one proposition must be true (although a proposition could be that nothing is happening).

Information about belief in a hypothesis A_j is contained in what is called the *evidential interval*. In order to define this, we must first define the *support*, $Spt(A_j)$, which is given by

$$Spt(A_j) = \sum_{A_i \subseteq A_j} m_A(A_i) \quad (C.3)$$

that is, the support for a hypothesis A_j is the sum of the masses of all propositions that are subsets of A_j (including A_j itself). The evidential interval is then easiest illustrated by:



Here, $\text{Pls}(A_j)$ represents the *plausibility* of A_j , that is, the degree to which the evidence fails to support its negation, and the difference between support and plausibility represents the residual ignorance, or uncertainty, $U_{\theta}(A_j)$. This concept is usually represented by $[\text{Spt}(A_j), \text{Pls}(A_j)]$, where actual numerical values are used within the square brackets.

The assignment of values for the masses $m_A(A_j)$ is, of course, problem dependent (and, in many cases, rather arbitrary), and may be time-dependent. However, once they have been assigned, masses for different times, knowledge sources or frames of discernment may be combined according to simple and credible rules to allow the evidential intervals for various propositions to be computed in a way that encompasses all of the evidence available. These rules will not be specified in detail in this report; instead, a very simple example will be worked through numerically.

Suppose information is sought from two completely independent knowledge sources (i.e. informants), on whether a particular vessel is friendly or unfriendly. The first source states that 50% of the evidence points to it being friendly, 20% points to it being unfriendly, and the remaining 30% could be interpreted either way. The second source gives estimates of 40%, 40% and 20% respectively. One would probably say that these data are largely consistent between informants, but for each informant contain significant self-contradiction and uncertainty. The question is: can we improve our knowledge by combining the data? The answer lies in Dempster's rule of combination, which is illustrated by the diagram below.

$m_A^1\{F\}$	0.5	$\{F\}$	ϕ	$\{F\}$
$m_A^1\{U\}$	0.2	ϕ	$\{U\}$	$\{U\}$
$m_A^1\{F, U\}$	0.3	$\{F\}$	$\{U\}$	$\{F, U\}$
		0.4	0.4	0.2
		$m_A^2\{F\}$	$m_A^2\{U\}$	$m_A^2\{F, U\}$

The possible objective situations, friendly or unfriendly, may be specified by a two-member set $\theta_A = \{F, U\}$. Since the beliefs of each informant can be expressed as divisions of a unit line (the vertical and horizontal axes respectively), it seems reasonable to express the combined beliefs of both informants as divisions of a unit square, as shown. We consider now a single division, the top right-hand corner of the square. This represents a measure of our combined belief that has been assigned by one informant to $\{F\}$, or friendly, and by the other to $\{F, U\}$, or indeterminate (the description *vacuous* is commonly applied). The only proposition to which we could reasonably assign this portion of our belief is the widest proposition consistent with both of these subsets, that is the intersection of the two subsets, $\{F\} \cap \{F, U\} = \{F\}$. This is indicated by the set description superimposed on the division.

The same sort of argument can be applied to all the divisions, as shown. We now have the problem that some of the belief is assigned to the empty set ϕ - this represents completely contradictory evidence to which we should assign no mass. The problem is overcome by assigning to each proposition in the combined evidence a mass that represents the ratio of the areas assigned to it and to all the non-empty sets. This is a simple instance of Dempster's rule, for which the general case is:

$$m_A^3(A_i) = (1 - k)^{-1} \sum_{\{i,j|A_i \cap A_j = A_i\}} m_A^1(A_i) m_A^2(A_j) \quad (C.4)$$

where

$$k = \sum_{\{i,j|A_i \cap A_j = \emptyset\}} m_A^1(A_i) m_A^2(A_j), \quad k \neq 1 \quad (C.5)$$

We can ignore the evidential interval assigned in all cases to the vacuous proposition $\{F, U\}$, which unsurprisingly turns out to be $[1, 1]$ (representing certainty that the vessel is either friendly or unfriendly), and for this simple binary choice the interval for $\{U\}$ can be deduced from that for $\{F\}$. A little simple arithmetic will assure us that the first informant implied an evidential interval for $\{F\}$ of $[0.5, 0.8]$, the second $[0.4, 0.6]$ and the combined evidence approximately $[0.58, 0.67]$. Have we improved our knowledge? The most obvious result is that the residual uncertainty U_0 has been reduced, from 0.3 for the first informant and 0.2 for the second, to 0.083; the second result is that the masses of evidence for the two elementary propositions are still not greatly different from each other. In other words, the combination has highlighted the essential nature of the evidence (that it is largely self-contradictory), whilst reducing the uncertainty.

If we now follow through the same calculation, but using as estimates for the beliefs of the two informants the sets $\{70\%, 10\%, 20\% \}$ and $\{80\%, 10\%, 10\% \}$, we find that the evidential interval for $\{F\}$ has changed from the two individual estimates of $[0.7, 0.9]$ and $[0.8, 0.9]$ to approximately $[0.93, 0.95]$. Here, two fairly high estimates for the likelihood of the vessel being friendly produce a very high combined likelihood, with little uncertainty. One should, however, be sure that the estimates are *independent*. If the problem had been, say, in weather forecasting, where the forecasters used the same data (and probably learned the same rules for manipulating them), such a combinational rule would not be appropriate.

We would expect the same sort of behaviour, for the case of independent evidence, with less-trivial propositions and/or more informants, conditions where a casual examination of the evidence would be much less likely to give a useful summary. It should be pointed out here that Dempster's rule is both commutative and associative, so evidence from an arbitrary number of sources may be combined in any order to give the same result.

In a similar manner to the combination technique described above, plausible rules have been developed for a number of operations on data sets. These include:

- ♦ *Translation* - the movement of information between different contexts, or frames of discernment. An example might be where one frame refers to vessels by type or class, and the other by properties such as displacement or magnetic signature. This operation requires some form of compatibility mapping or matrix, to allow one data set to be transformed to another frame for combination. The form of the translational rules depends on whether the frames of discernment are essentially independent, as in the suggestion above, or are different subsets of the same frame, such as vessel type and vessel class. When the different frames represent simply the same body of evidence at different times, this process is referred to as *projection*.


- ♦ *Discounting* - the modification of a mass distribution to take account of the reliability of a source. As an example, if the source is considered only 50% reliable, all masses are halved except that referring to the vacuous proposition (supporting all conclusions), which is expanded to maintain the unit sum. This allows information from sources of disparate reliability to be combined in a meaningful way.
- ♦ *Summarisation* - simplification of a body of evidence by eliminating those propositions for which the assigned mass is low.
- ♦ *Interpretation* - combination of the masses of evidence for and against a proposition, to provide a measure of its truthfulness.
- ♦ *Gisting* - determination of the proposition that best illustrates the general trend of a body of evidence. To obtain the gist, one first selects all the propositions with the equal greatest support. If there is only one, this is the gist. If there is more than one, the selection is narrowed to those with the lowest cardinality (number of elementary propositions). The gist is the remaining proposition (if there is only one), or the union of the remaining propositions.

In summary, evidential reasoning, which is a super-set of Dempster-Shafer theory, is a formal structure for reasoning about available information that may be incomplete and/or self-contradictory. It operates by combining all the information according to a credible set of rules, which require estimates to be made of the reliability of the information source, and the degree of belief that the source ascribes to each possible proposition.

Appendix D: Fuzzy Representations

D.1 Fuzzy Sets

In the very brief mention of fuzzy logic, above (Section 3.4), it was stated that an object either belonged in a given crisp set, or it didn't, but that it could have an intermediate membership in a fuzzy set. This is illustrated below, where $\chi_{\text{windy}}(v)$ represents the membership of a given wind speed v in the crisp or fuzzy set *windy*.



All this tells us is that, for the crisp set, all wind speeds below about 10 m/s (the actual figure is unimportant) would be considered not windy, or calm, whilst all speeds above this would be considered windy. (The membership at the most important speed, 10 m/s, is not defined.) In the fuzzy set, all speeds above 20 m/s are considered windy, but any below this have partial memberships in both the windy and calm sets - there is no uncomfortable ambiguity at any speed.

Zadeh (1973) recognises three types of operation on or between fuzzy sets; these are

- ♦ *negation* - characterised by the operator **not**,
- ♦ *connection* - characterised by operators like **and** and **or**, and
- ♦ *hedging* - characterised by operators like **very** and **quite**.

The first two types are familiar from propositional logic (Section 2.1 above) and crisp set theory. In fuzzy logic, as in propositional logic and crisp set theory, they have meanings consistent with common non-mathematical usage, and these meanings tend to the crisp-set meanings as the membership function χ tends to a step function. Hedges, on the other hand, represent somewhat arbitrary functions that tend to do-

nothing operations as the membership function tends to a step function, and modify membership functions in a way that is broadly consistent with their non-mathematical meanings. Combinations of all three types of operation give rise to composite functions, or *linguistic variables*, that are derived from the original membership functions, so that if, for example, the function *windy* has been defined, and *calm* is defined as *not windy*, then there is a precisely defined membership function for *not calm but not very windy*, whose values relate to those for *windy* in a commonsense way.

As with evidential reasoning above, this report will present in detail only a small subset of the available fuzzy-set operations, in order to demonstrate that their operation is reasonable when considered in conjunction with crisp-set theory and common usage. Let us consider first the negation operator. In a crisp set, this changes membership ($\chi = 1$) to non-membership ($\chi = 0$) - the equivalent with fuzzy sets is complementation (in the arithmetic sense) so that, with the above definitions,

$$\forall v \in V, \quad \chi_{\text{calm}}(v) \equiv 1 - \chi_{\text{windy}}(v) \quad (\text{D.1})$$

where V is the domain of v (which here would be the non-negative real numbers). In the notation usually employed for fuzzy-set theory (Kandel and Schneider, 1989), this would appear as:

$$\bar{W} \triangleq \neg W \triangleq \int_V (1 - \chi_W(v)) / v \quad (\text{D.2})$$

where W is used for the set *windy*. However, at the level considered in this report, this specialised notation will not be necessary.

Set negation is illustrated in the diagram below, where axis labels have been omitted for simplicity. The definition of Eqn. (D.1) makes sense from two points of view; if applied to crisp set, it produces the right answer, and the more a wind speed belongs to the set *windy*, the less it belongs to the set *not windy*.


In crisp set theory, the connective **and** represents set intersection, that is, if an object belongs to sets A and B it must belong to $A \cap B$. It follows that an object cannot belong to both A and $\neg A$ (**not** A), since the intersection of a set with its negation is an empty set. In terms of membership function χ , the crisp-set membership of A and B is 1 only if the memberships in A and B separately are both 1. In fuzzy-set theory, the equivalent operation is minimisation, that is:

$$\forall x \in X, \quad \chi_{A \text{ and } B}(x) \equiv \min(\chi_A(x), \chi_B(x)) \quad (\text{D.3})$$

Again, the definition appears reasonable; it tends to the crisp-set definition as the membership function approaches a step function, and it ensures that the degree to which an object belongs in the intersection of two sets cannot be greater than the

degree to which it belongs to either set. The equivalent for the or operator is, for similar reasons, maximisation.

We now have the situation that an object will generally have a non-zero membership of both a fuzzy set and its negation. If we consider this in relation to the linguistic variable *windy* and *not windy*, this appears to make little sense. However, again defining *calm* as *not windy*, this connective operation may be re-stated as any of *calm* and *windy*, *calm* and *not calm* or *not calm* and *not windy*, the last making sense in common usage. The diagram below shows how this function is derived from *windy* - it is, quite reasonably, a function that has low values everywhere except near the cross-over point between *calm* and *windy*.



In discussing the meanings assigned to hedges like *very*, it is perhaps clearer to start from common usage. There is, of course, no accepted quantitative definition of *very*, but we would expect to find three relationships between set membership for the sets *windy* and *very windy*:

- ♦ whatever the wind speed, its membership of *very windy* would be less than its membership of *windy*,
- ♦ for high wind speeds, both memberships would approach unity, and
- ♦ as the wind speed approaches zero, the ratio of memberships in *very windy* and *windy* would decrease.

For compatibility with crisp-set theory, the operation should have no effect on a step function of unit height. There are many functions that could be applied to χ_{windy} to produce $\chi_{\text{very windy}}$ and satisfy these requirements, and that commonly used is *concentration*, which is simply the squaring of the membership function, i.e.

$$\forall v \in V, \quad \chi_{\text{very windy}}(v) \equiv \chi_{\text{windy}}^2(v) \quad (\text{D.4})$$

This operation is the last of the examples illustrated in the figure above.

Fuzzy set theory, as detailed above, is essentially a descriptive tool. However, it can be used in *typicality theory* to derive an expectation for the magnitude, or range of magnitudes, for a variable, given incomplete information about its population. For example, given reasonable interpretations of the salient words, it is possible to define an expectation for wind speed from a statement such as "usually, the wind speed is between 2 and 15 m/s, but for about 10% of the time it is higher, and for almost 5% of the time it is lower". Perhaps more importantly, it is possible to check if given data on wind speed are consistent with such a statement, for example in a rule stating *if the wind speed is not typical then*. This is an example of fuzzy-set theory in the interpretation of fuzzy conditional statements, as briefly described below.

D.2 Fuzzy Logic

As anticipated in Section 3.4 above, a significant reason for investigating fuzzy logic is to allow for fuzzy inputs and outputs to condition-action clauses, e.g. "if the weather is *very windy* then sonar detection becomes *quite inefficient*". In order to achieve this, one must first cast the statement in a set-theoretical form. We will do this first using crisp-set theory, but following the example of Zadeh (1973), i.e.

$$\text{If } A \text{ then } B \text{ else } C \triangleq A \times B + (\neg A \times C) \quad (\text{D.5})$$

where A represents a member of a set of causes, and B and C are members of a set of consequences. Here, $A \times B$ represents the Cartesian product of the sets A and B that is, the set whose domain is all possible ordered pairs of members of A and B , and whose membership is **true** ($\chi = 1$) if and only if both members of the pair are members of their respective sets. The symbol $+$ here represents set union, which is possible here because the domains of both expressions that it joins are the same. Then any member of this domain represents a possible combination of cause and effect, and the expression evaluates to **true** if that effect is a consequence of that cause. Since we already have a representation of the union of two fuzzy sets (which is the same as $A \text{ or } B$), it only remains to define the Cartesian product of two fuzzy sets. By analogy with the definition of set intersection ($A \text{ and } B$), this is the fuzzy set in the combined domain whose membership is the minimum of the membership values of the member pair in their respective domains, i.e.

$$\forall \{(x, y) | x \in A, y \in B\}, \quad \chi_{A \times B}(x, y) = \min(\chi_A(x), \chi_B(y)) \quad (D.6)$$

When Equ. (D.5) is applied to a fuzzy system, any effect can now be associated with any cause by a membership value in the interval $[0,1]$, which can be considered, rather loosely, to be the probability of that effect given that cause. Where the sets of cause and effect have limited discrete domains, this is often represented by a matrix, referred to as a *fuzzy relational matrix*. Thus, given a numeric input for the cause, one may use the fuzzy relationship of Equ. (D.5) to obtain a fuzzy set representing the possible effects.

When the input to a condition-action clause is itself a fuzzy set, a further stage of processing is required, somewhat analogously to the convolution of an input signal with an impulse response to give an output signal. This process, known as the *compositional rule of inference*, makes use of the *max-min product*, defined by:


$$\chi_B(y) = \max_A(\min(\chi_A(x), \chi_{A \times B}(x, y))) \quad (D.7)$$

where $\chi_{A \times B}$ represents the fuzzy relational matrix referred to above. Thus, a fuzzy input x gives rise to a fuzzy output y , which may be interpreted, or passed on to a further decision-making step for similar treatment. The interpretation of the final fuzzy output will depend on whether the result is required as a binary decision (yes or no), or as the relative merits of a number of possible conclusions. In the latter case, the set memberships of the final set represent simply the required output. In the former case, it would be usual simply to make a decision based on which value of χ is the highest.

The final question to be answered is: why use fuzzy sets at all? The most obvious answer is for conditions where some or all of the inputs are inherently fuzzy, as in the example quoted in section 3.4. There is, however, another set of conditions where fuzzy logic may be used. This is where a system has significant negative feedback, so that approximate solutions will suffice, and at the same time the exact solutions to some of the decisions may be known in principle, but are too difficult to compute in real time. Such an application may not, when first encountered, seem to be justified, since there is no real evidence that the process described above will lead to a stable system. However, in many instances, e.g. steering a pilotless vehicle along a specified course, the rules may be obtained by consulting with experts who have, through long experience, discovered how to maintain a stable and accurate course, so there is a sort of *de facto* stability built into the rules.

Appendix E: Artificial Neural Networks

An artificial neural network is essentially a means of establishing a desired non-linear relationship between a number of input values and one or more output values. In order to achieve this, it is necessary to have a defined architecture containing processing elements, or *neurons*, and a means of training the system to produce the required output when supplied with inputs for which the correct output is known (Rumelhart and McClelland, 1987). Probably the best known such network is the Rumelhart back-propagation architecture, which is shown schematically below.



We consider a general strictly-layered three-layer network of this type, where the processing elements are indexed by k in the input layer, j in the hidden layer, and i in the output layer. We also define the output of a processing element as S_i if it is an output neuron, and s_j if it is a hidden neuron. A synaptic coupling (also referred to as a connection weight) between a hidden neuron and an output neuron is given by W_{ij} . Similarly, a coupling between an input node and a hidden neuron is given by w_{jk} . Finally, the threshold potential (also referred to as a bias) for an output neuron is given by V_i , and for a hidden neuron by v_j . The equation of state defining the network is (Muller and Reinhardt, 1990):

$$S_i = f(h_i)$$

where

$$h_i = \sum_j (W_{ij}s_j) - V_i \quad (\text{E.1})$$

with a similar expression for s_j . The non-linear transfer functions (also known as activation functions), $f(\cdot)$, are required to be continuous, differentiable and monotonically increasing, examples of which include:

$$f(h_i) = (1 + \exp(-\beta h_i))^{-1}$$

and

$$f(h_i) = \tanh(\beta h_i) \quad (\text{E.2})$$

During the training phase, the network iteratively adjusts weights, w , and bias values, v , to minimise the error function, D , between target values, ζ , and output values, $f(\cdot)$, for all classes, where

$$D(W_{i,j}, V_i, w_{j,k}, v_j) = \frac{1}{2} \sum_{\mu} \sum_i (\zeta_i^{\mu} - f(h_i^{\mu}))^2 \quad (\text{E.3})$$

In the Rumelhart back-propagation model (the most common and well established), parameter adjustment, such as $w_{n+1} = w_n + \delta w_n$ in the case of the connection weights, is generally achieved by application of a gradient descent procedure of the general form $w_{n+1} = w_n - \epsilon \nabla D(w_n)$. In the output layer, the parameter increment is proportional to the magnitude and direction of the derivative of the error function, and takes the form

$$\delta W_{i,j} = -\epsilon \frac{\partial D}{\partial W_{i,j}} = \epsilon \sum_{\mu} (\zeta_i^{\mu} - f(h_i^{\mu})) f'(h_i^{\mu}) \frac{\partial h_i^{\mu}}{\partial W_{i,j}} \quad (\text{E.4})$$

with similar rules for V_i , $w_{j,k}$ and v_j .

The error minimisation process thus involves the propagation of the output error deviation backward through the network. Details relating to the numerical implementation of back-propagation networks, including advice on scaling factors and extensions to the basic approach, can be found in Rumelhart and McClelland (1987), Muller and Reinhardt (1990), and Simpson (1990).

The principal application of artificial neural networks in AI is for the representation of objective causal relationships between a number of input values and one or more output values. Clearly, this would not be appropriate in cases where the relationship is well-known and simply calculable in real time. However, this is not always the case; it may take hours of computer time to calculate a known relationship, or the relationship may be known only as the result of experimental observations. In such cases, the parameters of a network may be trained to match a representative sample of calculations or observations, and the network then used as a simulation of the relationships. One subset of such relationships that is of particular interest is classification, where each output may represent a particular conclusion from the inputs, and all outputs go to zero except this, which goes to its maximum value. An

example of this type of classification is the deduction of ship type by underwater weapons on the basis of measured signatures.

One reservation that must always be borne in mind is that artificial neural networks do not contain intelligence in the sense of being based on models of the process involved - extrapolation outside the training area, or small pockets of aberrant behaviour not sampled within the area, can lead to wrong conclusions. However, they are extremely useful for multi-dimensional interpolation, and for limited extrapolation of well-behaved functions.

Appendix F: Glossary of Terms Used in this Report

The following consists of a set of terms from this report, each accompanied by a non-technical description (rather than an exact definition). References to other descriptions within the glossary are in *bold italic*. This Glossary has been conceptually clustered according to functional relationship.

Fundamental AI

Axiomise	Represent as a set of sentences in first order logic.
Combinational	Relating to all possible combinations of relevant information.
Condition-action statement	A logical rule of the form "If this <i>condition</i> holds, then this <i>action</i> is appropriate, else the other <i>action</i> is appropriate."
Conditional statement	Same as <i>condition-action statement</i> .
Connective	Used to describe how information should be combined in a <i>condition-action statement</i> . Usually <i>and</i> , <i>or</i> or <i>not</i> .
Disjunction	A series of propositions, one of which must be true in order for the overall proposition to be true. Also applied to the <i>or</i> connective.
Domain specific	Relating to a particular problem or class of problem.
If-then-else rule	Same as <i>condition-action statement</i> .
Inference	The use of rules of logic to draw conclusions from given information. The process of deriving new facts from old facts.
Logical inference	The drawing of conclusions where the rules of the problem and all input data are accurately known.
Null hypothesis	The default hypothesis against which other hypotheses are tested.
Parametric observation	An observation corresponding to a physical measurement.

Predicate	A statement about an object that, when applied to a specific argument, has a value of <i>True</i> or <i>False</i> .
Predicate calculus	An extension of <i>propositional logic</i> that allows for the description of objects that make up a proposition, and for reasoning about both the objects and the propositions.
Production rule	One of the rules in a <i>production system</i> , typically in the form of a <i>condition-action statement</i> .
Production system	A production system is a program which includes a body of knowledge (knowledge base) and an <i>inference engine</i> .
Propositional logic	Classical logic, which assumes that the problem solver has full information about input conditions and about the rules for handling information.
Quantifier	In <i>predicate calculus</i> , an operator such as <i>for all</i> , or <i>there exists</i> , used for making general statements about elements of a set.
Record	A set of related information that can be treated either as a single entity or as a number of <i>fields</i> .
Resolution	The underlying search and inference strategy of logic systems. Resolution is used to determine the truth of an assertion in logic systems free from contradictions.
Semantic network	A graphical representation of facts and relationships between them.
State transition network	Similar to a <i>semantic network</i> , but with emphasis on transitions rather than relationships between states.

Artificial Neural Networks

Activation function	A non-linear relationship between the weighted sum of inputs to a <i>neuron</i> , and the output from the neuron.
Artificial neural network	A computer simulation of loosely based on the brain which consists of at least one neuron and synapses. The neuron has a activation level and a transfer function. The synapses are the connection points for the neurons, and are made up of an input, a connection weight, and an output. The neurons may be connected in a complex network and they work in parallel with each other.

Back propagation	One of the most common techniques for training an <i>artificial neural network</i> .
Bias	A constant factor added to the weighted inputs to a <i>neuron</i> .
Neural network	Loose description of <i>artificial neural network</i> .
Neuron	The processing element that takes a number of inputs, together with <i>weights</i> and a <i>bias</i> , and produces an output that reflects the manner in which the inputs react.
Rumelhart architecture	One of the most common dispositions for an <i>artificial neural network</i> .
Synaptic coupling	A coupling, with its associated <i>weight</i> , that represents the influence one input, or intermediate combination of inputs, has on a following <i>neuron</i> .
Threshold potential	An alternative term for <i>bias</i> .
Transfer function	The same as <i>activation function</i> .
Weight	A factor quantifying the importance of a <i>synaptic coupling</i> .

Quantitative Inference

<i>A priori</i>	<i>A priori</i> information is that available for a given situation before an attempt to derive conclusions.
<i>A posteriori</i>	<i>A posteriori</i> information is that resulting from the drawing of conclusions about a given situation.
Bayesian	Bayesian theory is a means of drawing inferences from probability distributions.
Body of evidence	In <i>evidential reasoning</i> , the information that leads to the assignment of <i>masses</i> for <i>propositional statements</i> .
Cardinality	In <i>crisp</i> or <i>fuzzy sets</i> with discrete domains, the number of domain items belonging to the set.
Cartesian product	In <i>crisp</i> or <i>fuzzy sets</i> , the set for which each element represents one possible combination of elements from each of two or more component sets, and all possible combinations are allowable.

Compositional rule of inference	In <i>fuzzy logic</i> , a rule computing how a fuzzy or <i>crisp</i> value for a variable can lead to a fuzzy value for a new variable.
Crisp	A crisp value is one that has a single numeric value, and is therefore not <i>fuzzy</i> .
Dempster-Shafer theory	A means of handling uncertain and/or incomplete data that is the basis of <i>evidential reasoning</i> .
Discounting	In <i>evidential reasoning</i> , a means of allowing for the relative reliability of inconsistent data.
Evidential interval	In <i>evidential reasoning</i> , the division of confidence in a hypothesis into <i>support</i> , <i>plausibility</i> and (implicitly) <i>uncertainty</i> .
Evidential reasoning	A body of techniques for automated reasoning from evidence that may be uncertain and/or incomplete.
Focal element	In <i>evidential reasoning</i> , a hypothesis with non-zero <i>mass</i> .
Frame	<ol style="list-style-type: none"> 1. A group of information about particular objects or events. 2. Also, in <i>evidential reasoning</i>, the same as <i>frame of discernment</i>.
Frame of discernment	The set of all possible values for a variable in its domain, particularly in <i>evidential reasoning</i> .
Fusion	The combination of data from different sources.
Fuzzy logic	A logical system in which some or all values encountered are <i>fuzzy values</i> , and in which the <i>condition-action statements</i> may be in an approximate form.
Fuzzy relational matrix	A <i>fuzzy set</i> whose domain is the <i>Cartesian product</i> of the domains of two or more component sets, not necessarily fuzzy, and whose <i>membership function</i> represents the membership of the component domain elements in a set describing a given relationship.
Fuzzy set	A set for which each element, or position in the domain, can have a partial membership, that is, can have a specific degree of membership, in the range 0 to 1, where 0 represents non-membership, or <i>False</i> , and 1 represents membership, or <i>True</i> . The value of the partial membership is known as the <i>membership function</i> .
Fuzzy value	A value represented by a <i>fuzzy set</i> over its domain.

Gisting	In <i>evidential reasoning</i> , a technique for finding the most <i>pointed</i> hypothesis, that is, that hypothesis from those with equal maximum <i>support</i> which has the fewest component <i>predicates</i> .
Hedge	In <i>fuzzy logic</i> , an operator on a <i>fuzzy set</i> that modifies it in a particular, and usually arbitrary, fashion. Typical hedges would be <i>very</i> and <i>approximately</i> .
Insufficient reason	The rule of insufficient reason states that, if there is no reason to believe that the probabilities of a given set of hypotheses are different, the hypotheses should be assigned equal probability.
Interpretation	In <i>evidential reasoning</i> , the combination of evidence for and against a proposition to provide a measure of confidence in its truthfulness.
Likelihood ratio	The ratio of probabilities for and against a hypothesis.
Linguistic variable	A <i>fuzzy value</i> that is expressed in (possibly constrained) natural language, such as <i>fairly long</i> .
Mass	In <i>evidential reasoning</i> , the value of belief assigned to a given basic <i>predicate</i> . The masses for all basic predicates within the <i>frame of discernment</i> sum to unity.
Maximum relative entropy	A means of updating a probability distribution involving the minimisation of the information required to be considered.
Max-min product	In <i>fuzzy set</i> theory, a means of deriving a fuzzy set from an input fuzzy set and a rule represented by a <i>fuzzy relational matrix</i> .
Membership function	In <i>fuzzy set</i> theory, the degree to which a position in the domain belongs to a particular set.
Plausibility	In <i>evidential reasoning</i> , the plausibility of a hypothesis is the sum of masses not assigned to its negation, that is, the degree to which the evidence fails to refute the hypothesis.
Power set	The power set of a given set θ of hypotheses is the set of sets whose domain is all possible subsets of θ , including θ itself. Usually denoted by 2^θ .
Probabilistic inference	Application of <i>Bayesian</i> theory.

Probabilistic logic	A logical system in which reasoning is carried out using separate conceptual worlds for all allowable combinations of conditions.
Projection	In <i>evidential reasoning</i> , the movement of information through different contexts representing discrete times, to allow for the simulation of time-dependent systems.
Propositional statement	In <i>evidential reasoning</i> , a statement that is the union of one or more basic statements.
Statistical inference	The use of statistics from past experience to calculate probabilities of future events.
Summarisation	In <i>evidential reasoning</i> , the simplification of a <i>body of evidence</i> by eliminating those <i>propositional statements</i> for which the assigned <i>mass</i> is low
Support	In <i>evidential reasoning</i> , the sum of the <i>masses</i> assigned to the basic <i>predicates</i> that constitute a <i>propositional statement</i> . It is a measure of confidence in that statement.
Translation	In <i>evidential reasoning</i> , the movement of information between <i>frames of discernment</i> .
Typicality theory	In <i>fuzzy logic</i> , the establishment of fuzzy criteria to summarise the expectations for objects or events.
Uncertain reasoning	Any form of logic that can deal with information that is approximate, missing or contradictory, or where the relationships between objects and events are not known exactly.
Uncertainty	In <i>evidential reasoning</i> , the difference between <i>support</i> (confidence in an event) and <i>plausibility</i> (lack of confidence in its negation).
Universe of discourse	The domain of a problem, that is, the aggregate of objects and events that are relevant to its solution. Also used in <i>evidential reasoning</i> as equivalent to <i>frame of discernment</i> .
Vacuous proposition	In <i>evidential reasoning</i> , the proposition that contains all possible <i>predicates</i> , and is therefore intrinsically true.
World model	A conceptual description of all events and objects relevant to a particular problem. Different from a <i>universe of discourse</i> in that it assigns values or states to all objects and events.

World picture A loose expression that can mean either *universe of discourse* or *world model*.

Decision Updating

Nonmonotonic reasoning A method of reasoning depending on making the best estimate of conditions at any decision point, with a facility for back-tracking if the assumptions lead to an inconsistent conclusion.

Classification

Attribute vector A number of *parametric observations* represented as a vector in a multi-dimensional Cartesian coordinate system.

Classification Sorting of events or objects into different categories, usually according to objectively measurable criteria.

Cluster analysis *Classification* according to similarity of attributes, e.g. according to a *resemblance matrix*.

Figure of merit A means of *classification* of events or objects according to conformity with given *attribute vectors*.

Resemblance matrix In *cluster analysis*, a matrix used to describe the resemblance between any two of a set of objects or events, based on their positions in a non-dimensional parameter space.

Templating *Classification* by assessing the compliance of an object or event with a number of independent qualitative or quantitative criteria.

Application of Artificial Intelligence to Decision Making
in Mine Countermeasures

T.M. Mansell, D.R. Skinner and K.K. Benke

(DSTO-TR-0279)

DISTRIBUTION LIST

DEFENCE ORGANISATION

Defence Science and Technology Organisation

Chief Defence Scientist	} shared copy
FAS Science Policy	
AS Science Corporate Management	
Counsellor Defence Science, London (Doc Data Sheet only)	
Counsellor Defence Science, Washington (Doc Data Sheet only)	
Scientific Adviser to Thailand MRD (Doc Data Sheet Only)	
Scientific Adviser to the DRC (Kuala Lumpur) (Doc Data Sheet Only)	
Senior Defence Scientific Adviser/Scientific Adviser Policy and Command (shared copy)	
Navy Scientific Adviser	
Scientific Adviser - Army (Doc Data Sheet only)	
Air Force Scientific Adviser	
Director Trials	

Aeronautical and Maritime Research Laboratory

Director, AMRL
Chief, Maritime Operations Division
Dr D.D. Richardson
T.M. Mansell
D.R. Skinner
K.K. Benke

DSTO Library

Library Fishermens Bend
Library Maribyrnong
Main Library DSTOS (2 copies)
Library, MOD, Pyrmont (Doc Data Sheet only)

Defence Central

OIC TRS, Defence Central Library
Officer in Charge, Document Exchange Centre, 1 copy
DEC requires, as permitted by the release limitations, the following copies to meet
exchange agreements under their management:
*US Defence Technical Information Centre, 2 copies
*UK Defence Research Information Centre, 2 copies
*Canada Defence Scientific Information Service, 1 copy
*NZ Defence Information Centre, 1 copy
National Library of Australia, 1 copy
The Director, Qinghua University, Beijing, 1 copy
The Librarian, Beijing University of Aeronautics, 1 copy
The Librarian, Institute of Mechanics, Chinese Academy of Sciences, 1 copy

DISTRIBUTION LIST
(Continued)

(DSTO-TR-0279)

Defence Intelligence Organisation
Library, Defence Signals Directorate (Doc Data Sheet only)

Army

Director General Force Development (Land) (Doc Data Sheet only)
ABCA Office, G-1-34, Russell Offices, Canberra (4 copies)
NAPOC QWG Engineer NBCD c/- DENGSR-A, HQ Engineer Centre

Navy

ASSTASS, APW2-1-OA2, Anzac Park West, Canberra
COMAUSMINFOR (CMDR R. Baker), HMAS Waterhen, Waverton NSW 2060
DASD (CMDR C.W. Donald), Anzac Park West Offices APW2-4-01, Canberra ACT 2600
DCFD(SEA) (CAPT G.R. Walpole), Russell Offices B4-4-04, Canberra ACT 2600
DDMWD (CMDR G.J. Kelly), Russell Offices B-4-01, Canberra ACT 2600
DNCSE (CAPT P. Maguire), Campbell Park Offices CP1-6-12, Canberra ACT 2600
DNW (CAPT D.J. Ramsay), Russell Offices A-3-16, Canberra ACT 2600
DRANSWARS, HMAS Watson, Watson's Bay NSW 2030
FONTC (RADM P.D. Birggs), HMAS Cerberus, Crib Point Vic 3910
MHCPD (CAPT E.M. Asker), Campbell Park Offices CP-2-3-11, Canberra ACT 2600
MRARANRAU, Dept. of Defence, 54-56 Miller St, North Sydney NSW 2060
MWSCPD (CMD M.S. Welford), Campbell Park Offices CP2-2-19, Canberra ACT 2600

DSTO

Dr S.P. Goss, AOD, AMRL Fishermens Bend
Dr S.T. Hood, ITD, DSTO Salisbury
Dr R.F. Barrett, MOD, DSTO Salisbury
Dr D.F. Liebing, MOD, DSTO Salisbury

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy
Library
Head of Aerospace and Mechanical Engineering
Deakin University, Serials Section (M list)), Deakin University Library, Geelong, 3217,
Senior Librarian, Hargrave Library, Monash University

OTHER ORGANISATIONS

NASA (Canberra)
AGPS

ABSTRACTING AND INFORMATION ORGANISATIONS

INSPEC: Acquisitions Section Institution of Electrical Engineers
Library, Chemical Abstracts Reference Service
Engineering Societies Library, US
American Society for Metals
Documents Librarian, The Center for Research Libraries, US

DISTRIBUTION LIST
(Continued)

(DSTO-TR-0279)

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK
Library - Exchange Desk, National Institute of Standards and
Technology, US

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA				1. PAGE CLASSIFICATION UNCLASSIFIED	
				2. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
3. TITLE Application of artificial intelligence in decision making in mine countermeasures			4. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
5. AUTHOR(S) T.M. Mansell, D.R. Skinner and K.K. Benke			6. CORPORATE AUTHOR Aeronautical and Maritime Research Laboratory PO Box 4331 Melbourne Vic 3001		
7a. DSTO NUMBER DSTO-TR-0279		7b. AR NUMBER AR-009-442		7c. TYPE OF REPORT Technical Report	
8. DOCUMENT DATE December 1995					
9. FILE NUMBER G6/4/8-4517		10. TASK NUMBER NAV 92/315		11. TASK SPONSOR	
				12. NO. OF PAGES 57	
				13. NO. OF REFERENCES 32	
14. DOWNGRADING/DELIMITING INSTRUCTIONS Not applicable.				15. RELEASE AUTHORITY Chief, Ship Structures and Materials Division	
16. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT Approved for public release OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE CENTRE, DIS NETWORK OFFICE, DEPT OF DEFENCE, CAMPBELL PARK OFFICES, CANBERRA ACT 2600					
17. DELIBERATE ANNOUNCEMENT Announcement of this report is unlimited					
18. CASUAL ANNOUNCEMENT YES/NO (Cross out whichever is not applicable)					
19. DEFTEST DESCRIPTORS Artificial intelligence; Mine countermeasures; Neural networks; Decision making					
20. ABSTRACT A number of analytic techniques used in Artificial Intelligence are examined in the context of decision making in mine countermeasures. Attention is directed at five major techniques, involving statistical inference, probabilistic inference, evidential reasoning, fuzzy logic and artificial neural networks. In the cases of statistical inference and evidential reasoning, solutions to appropriate problems are described. Eleven other techniques are dealt with more briefly, in most cases with worked examples of appropriate naval application. The main conclusion reached is that, in view of the probable shortage of accurate information under operational conditions, evidential reasoning and fuzzy logic are likely to be the most appropriate means for presenting relevant data to decision makers, and that artificial neural networks will be useful for representing complicated or empirical relationships between observed factors.					